



EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent:
18.09.2013 Bulletin 2013/38

(51) Int Cl.:
H04N 1/60 (2006.01)

(21) Application number: **03722835.0**

(86) International application number:
PCT/GB2003/001948

(22) Date of filing: **02.05.2003**

(87) International publication number:
WO 2003/094505 (13.11.2003 Gazette 2003/46)

(54) **IMAGE REPRESENTATION METHOD AND APPARATUS**
VERFAHREN UND GERÄT ZUR BILDDARSTELLUNG
PROCEDE ET APPAREIL DE REPRESENTATION D'IMAGE

(84) Designated Contracting States:
DE ES FR GB

(30) Priority: **03.05.2002 GB 0210189**

(43) Date of publication of application:
06.04.2005 Bulletin 2005/14

(73) Proprietor: **Apple Inc.**
Cupertino, CA 95014 (US)

(72) Inventor: **Finlayson, Graham,**
School of Information Systems
Norwich NR4 7TJ (GB)

(74) Representative: **Lang, Johannes**
Bardehle Pagenberg Partnerschaft
Patentanwälte, Rechtsanwälte
Prinzregentenplatz 7
81675 München (DE)

(56) References cited:
EP-B1- 0 663 094 US-A- 5 341 439
US-A- 5 450 502

- **NERWIN D H ET AL: "Comparison of eight color and gray scales for displaying continuous 2D data" PROCEEDINGS OF THE HUMAN FACTORS AND ERGONOMICS SOCIETY 37TH ANNUAL MEETING. DESIGNING FOR DIVERSITY, PROCEEDINGS OF 37TH ANNUAL MEETING ON THE HUMAN FACTORS AND ERGONOMICS SOCIETY, SEATTLE, WA, USA, 11-15 OCT. 1993, pages 1330-1334 vol.2, XP008025256 1993, Santa Monica, CA, USA, Human Factors & Ergonomics Soc, USA**

- **HANNA T. HASKETT, ARUN K. SOOD: 'Aided/automatic target detection using reflexive hyperspectral imagery for airborne applications', [Online] December 1998, ARMY COMMUNICATIONS-ELECTRONICS COMMANDFORT BELVOIR VA NIGHT VISION AND ELECTRONICS SENSORS DIRECTORATE Retrieved from the Internet: <URL:http://www.dtic.mil/cgi-bin/GetTRDoc?A=D=ADA399070&Location=U2&doc=GetTRDoc.pdf> [retrieved on 2010-02-05]**
- **IOANNIS PITAS ET AL: "MULTIVARIATE ORDERING IN COLOR IMAGE FILTERING", 1 September 1991 (1991-09-01), IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, PAGE(S) 247 - 259, XP000232805, ISSN: 1051-8215**
- **WEEKS A R ET AL: "jistogram specification of 24-bit color images in the color difference (C-Y) color space", PROCEEDINGS OF THE SPIE - THE INTERNATIONAL SOCIETY FOR OPTICAL ENGINEERING, SPIE, US, vol. 3646, 1 January 1999 (1999-01-01), pages 319-329, XP002381195, ISSN: 0277-786X**
- **'Rankit', [Online] 20 April 2009, Wikipedia Retrieved from the Internet: <URL:http://en.wikipedia.org/wiki/Rankit> [retrieved on 2010-02-05]**
- **BHAT D N ET AL: "ORDINAL MEASURES FOR IMAGE CORRESPONDENCE", IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, IEEE SERVICE CENTER, LOS ALAMITOS, CA, US, vol. 20, no. 4, 1 April 1998 (1998-04-01), pages 415-423, XP000766682, ISSN: 0162-8828, DOI: 10.1109/34.677275**

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

- JIA-LIN CHEN ET AL: "Rotation and gray scale transform invariant texture identification using wavelet decomposition and hidden Markov model", IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE USA, vol. 16, no. 2, February 1994 (1994-02), pages 208-214, ISSN: 0162-8828
- PIETIKAINEN M ET AL: "Accurate color discrimination with classification based on feature distributions", PROCEEDINGS OF THE 13TH INTERNATIONAL CONFERENCE ON PATTERN RECOGNITION IEEE COMPUT. SOC. PRESS LOS ALAMITOS, CA, USA, vol. 3, 1996, pages 833-838 vol.3, ISBN: 0-8186-7282-X

Description

[0001] Colour can potentially provide useful information to a variety of computer vision tasks such as image segmentation, image retrieval, object recognition and tracking. However, for it to be helpful in practice, colour must relate directly to the intrinsic properties of the imaged objects and be independent of imaging conditions such as scene illumination and the imaging device. To this end many *invariant* colour representations have been proposed in the literature. Unfortunately, recent work [3] has shown that none of them provides good enough practical performance.

Summary of the Invention

[0002] The invention provides a method and an apparatus for representing, characterising, transforming or modifying a colour image as defined in the appended independent claims. Preferred or advantageous features of the invention are set out in dependent subclaims.

[0003] In this paper we propose a new colour invariant image representation based on an existing grey-scale image enhancement technique: histogram equalisation. We show that provided the rank ordering of sensor responses are preserved across a change in imaging conditions (lighting or device) a histogram equalisation of each channel of a colour image renders it invariant to these conditions. We set out theoretical conditions under which rank ordering of sensor responses is preserved and we present empirical evidence which demonstrates that rank ordering is maintained in practice for a wide range of illuminants and imaging devices. Finally we apply the method to an image indexing application and show that the method outperforms all previous invariant representations, giving close to perfect illumination invariance and very good performance across a change in device.

1 Introduction

[0004] It has long been argued that colour (or *RGB*) images provide useful information which can help in solving a wide range of computer vision problems. For example it has been demonstrated [14, 1] that characterising an image by the distribution of its colours (*RGBs*) is an effective way to locate images with similar content from amongst a diverse database of images. Or that a similar approach [19] can be used to locate objects in an image. Colour has also been found to be useful for tasks such as image segmentation [12, 13] and object tracking [7, 16]. Implicit in these applications of colour is the assumption that the colours recorded by devices are an inherent property of the imaged objects and thus a reliable cue to their identity. Unfortunately a careful examination of image formation reveals that this assumption is not valid. The *RGB* that a camera records is more properly a measure of the light reflected from the surface of an object and while this does depend in part on characteristics of the object, it depends in equal measure on the composition of the light which is incident on the object in the first place. So, an object that is lit by an illuminant which is itself reddish, will be recorded by a camera as more red than will the same object lit under a more bluish illuminant. That is, image *RGBs* are *illumination dependent*. In addition image colour also depends on the properties of the recording device. Importantly, different imaging devices have different sensors which implies that an object that produces a given *RGB* response in one camera might well produce a quite different response in a different device.

[0005] In recognition of this fact many researchers have sought modified image representations such that one or more of these dependencies are removed. Researchers have to-date, concentrated on accounting for illumination dependence and typically adopt one of two approaches: colour invariant [9, 5, 6, 10, 18] or colour constancy [11, 8] methods. Colour invariant approaches seek transformations of the image data such that the transformed data is illuminant independent whereas colour constancy approaches set out to determine an estimate of the light illuminating a scene and provide this estimate in some form to subsequent vision algorithms.

[0006] Colour constancy algorithms, in contrast to invariant approaches can deliver true object colours. Moreover, colour invariants can be calculated post-colour constancy processing but the converse is not true. This said, colour constancy has proven to be a harder problem to solve than colour invariants. Most importantly however, it has been demonstrated [8, 3] that the practical performance of neither approach is good enough to facilitate colour-based object recognition or image retrieval. Moreover, none of the methods even attempts to account for device dependence.

[0007] In this paper we seek to address the limitations of these existing approaches by defining a new image representation which we show is both illumination independent and (in many cases) also device independent. Our method is based on the observation that while a change in illumination or device leads in practice to significant changes in the recorded *RGBs*, the rank orderings of the responses of a given sensor are largely preserved. In fact, we show in this paper (§3) that under certain simplifying assumptions, invariance of rank ordering follows directly from the image formation equation. In addition we present an empirical study (§4) which reveals that the preservation of rank ordering holds in practice both across a wide range of illuminants and a variety of imaging devices. Thus, an image representation which is based on rank ordering of recorded *RGBs* rather than on the *RGBs* themselves offers the possibility of accounting for both illumination and device dependence at the same time.

[0008] To derive an image representation which depends only on rank orderings we borrow a tool which has long been used by the image processing community for a quite different purpose. The technique is histogram equalisation and is typically applied to grey-scale images to produce a new image which is enhanced in the sense that the image has more contrast and thus conveys more information. In some cases this results in a visually more pleasing image. But in a departure from traditional image processing practice, we apply the procedure not to a grey-scale image, but rather to each of the R , G , and B channels of a colour image independently of one another. We show that provided two images differ in such a way as to preserve the rank ordering of pixel values in each of the three channels then an application of histogram equalisation to each of the channels of the two images results in a pair of equivalent images. Thus provided a change in illuminant or device preserves rank ordering of pixel responses the application of histogram equalisation will provide us with an invariant representation of a scene which might subsequently be of use in a range of vision applications.

[0009] Of course the reader may be surprised that we propose something so simple: histogram equalisation is a common tool. Paradoxically however, histogram equalising R , G , and B channels of an image is generally discouraged because this results in unnatural pseudo-colours. For our purposes, however - recognition or tracking - pseudo-colours suffice. We demonstrate this (§5) by applying the method to the problem of colour indexing: we show that the method out performs all previous approaches and in the case of a change in illumination provides close to perfect indexing.

[0010] The scientific article by Pietikäinen et al. "Accurate color discrimination with classification based on feature distributions", ICPR'96, 1996, discloses a method for the retrieval of images based on the comparison of color image histograms. The illumination change invariance is achieved by a calibration procedure.

[0011] The scientific article by Chen et al. "Rotation and gray scale transform invariant texture identification using wavelet decomposition and hidden Markov model" in Transactions on Pattern Analysis and Machine intelligence, Vol. 16, February 1994, discloses a method for texture recognition where a gray scale transform invariance is achieved by an histogram equalization of the image intensity value.

2 Background

[0012] We base our work on a simple model of image formation in which the response of an imaging device to an object depends on three factors: the light by which the object is lit, the surface reflectance properties of the object, and the properties of the device's sensors. We assume that a scene is illuminated by a single light characterised by its spectral power distribution which we denote $E(\lambda)$ and which specifies how much energy the source emits at each wavelength (λ) of the electromagnetic spectrum. The reflectance properties of a surface are characterised by a function $S(\lambda)$ which defines what proportion of light incident upon it the surface reflects on a per- wavelength basis. Finally a sensor is characterised by $R_k(\lambda)$, its spectral sensitivity function which specifies its sensitivity to light energy at each wavelength of the spectrum. The subscript k denotes that this is the k th sensor. Its response is defined as:

$$p_k = \int_{\omega} E(\lambda) S(\lambda) R_k(\lambda) d\lambda, \quad k = 1, \dots, m \quad (1)$$

where the integral is taken over the range of wavelengths ω : the range for which the sensor has non- zero sensitivity. In what follows we assume that our devices (as most devices do) have three sensors ($m = 3$) so that the response of a device to a point in a scene is represented by a triplet of values: (p_1, p_2, p_3) . It is common to denote these triplets as R , G , and B or just RGB s and so we use the different notations interchangeably throughout. An image is thus a collection of RGB s representing the device's response to light from a range of positions in a scene.

[0013] Equation (1) makes clear the fact that a device response depends both on properties of the sensor (it depends on $R_k(\lambda)$) and also on the prevailing illumination on ($E(\lambda)$). That is, responses are both device and illumination dependent. It follows that if no account is taken of these dependencies, an RGB cannot correctly be considered to be an intrinsic property of an object and employing it as such is quite likely to result in poor results.

[0014] An examination of the literature reveals many attempts to deal with the illumination dependence problem. One approach is to apply a correction to the responses recorded by a device to account for the colour of the prevailing scene illumination. Provided an accurate estimate of the scene illumination can be obtained, such a correction accounts well for the illumination dependence, rendering responses colour constant: that is stable across a change in illumination. The difficulty with this approach is the fact that estimating the scene illuminant is non-trivial. In 1998 Funt *et al* [8] demonstrated that existing colour constancy algorithms are not sufficiently accurate to make such an approach viable. More recent work [11] has shown that for simple imaging conditions and given good device calibration the colour constancy approach can work.

[0015] A different approach is to derive from the image data some new representation of the image which is invariant

to illumination. Such approaches are classified as colour (or illuminant) invariant approaches and a wide variety of invariant features have been proposed in the literature. Accounting for a change in illumination colour is however difficult because, as is clear from Equation (1), the interaction between light, surface and sensor is complex. Researchers have attempted to reduce the complexity of the problem by adopting simple models of illumination change. One of the simplest models is the so called *diagonal* model in which it is proposed that sensor responses under a pair of illuminants are related by a diagonal matrix transform:

$$\begin{pmatrix} R^c \\ G^c \\ B^c \end{pmatrix} = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} R^o \\ G^o \\ B^o \end{pmatrix} \quad (2)$$

where the superscripts o and c characterise the pair of illuminants. The model is widely used, and has been shown to be well justified under many conditions [7]. Adopting such a model one simple illuminant invariant representation of an image can be derived by applying the following transform:

$$R' = \frac{R}{R_{ave}}, \quad G' = \frac{G}{G_{ave}}, \quad B' = \frac{B}{B_{ave}} \quad (3)$$

where the triplet $(R_{ave}, G_{ave}, B_{ave})$ denotes the mean of all *RGBs* in an image. It is easy to show that this so called *Greyworld* representation of an image is illumination invariant provided that Equation (2) holds.

[0016] Many other illuminant invariant representations have been derived, in some cases [10] by adopting different models of image formation. All derived invariants however share two common failings: first it has been demonstrated that when applied to the practical problem of image retrieval none of these invariants affords good enough performance across a change in illumination. Secondly, none of these approaches considers the issue of device invariance.

[0017] Device invariance occurs because different devices have different spectral sensitivity functions (different R_k in Equation (1)) but also because the colours recorded by a device are often not linearly related to scene radiance as Equation (1) suggests, but rather are some non-linear transform of this:

$$p_k = f \left(\int_{\omega} E(\lambda) S(\lambda) R_k(\lambda) d\lambda \right), \quad k = 1, \dots, m \quad (4)$$

[0018] The transform $f()$ is deliberately applied to *RGB* values recorded by a device for a number of reasons. First, many captured images will eventually be displayed on a monitor. Importantly colours displayed on a screen are not a linear function of the *RGBs* sent to the monitor. Rather, there exists a power function relationship between the incoming voltage and the displayed intensity. This relationship is known as the gamma of the monitor, where gamma describes the exponent of the power function [15]. To compensate for this gamma function images are usually stored in a way that reverses the effect of this transformation: that is by applying a power function with exponent of $1/\gamma$, where γ describes the gamma of the monitor, to the image *RGBs*. Importantly monitor gammas are not unique but can vary from system to system and so images from two different devices will not necessarily have the same gamma correction applied. In addition to gamma correction other more general non-linear "tone curve" corrections are often applied to images so as to change image contrast with the intention of creating a visually more pleasing image. Such transformations are device, and quite often, image dependent and so lead, inevitably to device dependent colour. In the next section we address the limitations of existing invariant approaches by introducing a new invariant representation.

3 Histogram Equalisation for Colour Invariance

[0019] Let us begin by considering again the diagonal model of image formation defined by Equation (2). We observe that one implication of this model of illumination change is that the rank ordering of sensor responses is preserved under a change of illumination. To see this, consider the responses to a single sensor R , such that R^o_i represents the response to a surface i under an illuminant o . Under a second illuminant, which we denote c , the surface will have response R^c_i and the pair of sensor responses are related by:

$$R_i^c = \alpha R_i^o \quad (5)$$

[0020] Equation (5) is true for all surfaces (that is, $\forall i$). Now, consider a pair of surfaces, i and j , viewed under illuminant o and suppose that $R_i^o > R_j^o$, then it follows from Equation (5) that:

$$R_i^o > R_j^o \Rightarrow \alpha R_i^o > \alpha R_j^o \Rightarrow R_i^c > R_j^c \quad \forall i, j, \quad \forall \alpha \quad (6)$$

[0021] That is, the rank ordering of sensor responses within a given channel is invariant to a change in illumination.

[0022] Thus, if what we seek is an image representation which is invariant to illumination we can obtain one by considering not the pixel values themselves but rather the relative ordering of these values. There are a number of ways we might employ this rank ordering information to derive an invariant representation, we set forth one such method here which we will demonstrate has a number of interesting properties. To understand our method consider a single channel of an *RGB* image recorded under an illuminant o where without loss of generality we restrict the range of R^o to be on some finite interval: $R^o \in [0 \dots R/\max]$. Now, consider further a value $R_j^o \in [0 \dots R/\max]$ where R_j^o is not necessarily the value of any pixel in the image. Let us define by $P(R^o < R_j^o)$, the number of pixels in an image with a value less than or equal to R_j^o . Under a second illuminant, c , a pixel value R^o under illuminant o is mapped to a corresponding value R_c . We denote by $P(R_c < R_j^c)$ the number of pixel values in the second image whose value is less than R_j^c . Assuming that the illumination change preserves rank ordering of pixels we have the following relation:

$$P(R^c < R_j^c) = P(R^o < R_j^o) \quad (7)$$

[0023] That is, the number of pixels in our image under illuminant o which have a value less than R_j^o is equal to the number of pixels in the image under illuminant c which have a value less than the transformed pixel value R_j^c ; a change in illumination preserves cumulative proportions. Given this, we define one channel of the invariant image representation thus:

$$R_i^{inv} = \frac{R_{max}}{N_{pix}} P(R^o \leq R_i^o) = P(R^c \leq R_i^c) \quad (8)$$

where N_{pix} is the number of pixels and the constant R_{max}/N_{pix} ensures that the invariant image has the same range of values as the input image. Repeating the procedure for each channel of a colour image results in the required invariant image.

[0024] The reader familiar with the image processing literature might recognise Equation (8). Indeed this transformation of image data is one of the simplest and most widely used methods for image enhancement and is commonly known as histogram equalisation. Histogram equalisation is an image enhancement technique originally developed for a single channel, or grey-scale, image. The aim is to increase the overall contrast in the image since doing so typically brightens dark areas of an image, increasing the detail in those regions which in turn can sometimes result in a more pleasing image. Histogram equalisation achieves this aim by transforming an image such that the histogram of the transformed image is as close as possible to a uniform histogram. The approach is justified on the grounds that amongst all possible histograms, a uniformly distributed histogram has maximum entropy. Maximising the entropy of a distribution maximises its information and thus histogram equalizing an image maximises the information content of the output image. Accepting the theory, to histogram equalise an image we must transform the image such that the resulting image histogram is uniform. Now, suppose that x_i represents a pixel value in the original image and x'_i its corresponding value in the transformed image. We would like to transform the original image such that the proportion of pixels less than x'_i in the transformed image is equal to the proportion of image pixels less than x_i in the original image, and that moreover the histogram of the output image is uniform. This implies:

$$\int_0^{x_i} p(x) dx = \int_0^{x'_i} q(x) dx = \frac{N_{pix}}{x_{max}} \int_0^{x'_i} dx \quad (9)$$

where $p(x)$ and $q(x)$ are the histograms of the original and transformed images respectively. Evaluating the right-hand integral we obtain:

$$\frac{x_{max}}{N_{pix}} \int_0^{x_i} p(x) dx = x'_i \quad (10)$$

[0025] Equation (10) tells us that to histogram equalise an image we transform pixel values such that a value x_i in the original image is replaced by the proportion of pixels in the original image which are less than or equal to x_i . A comparison of Equation (8) and Equation (10) reveals that, disregarding notation, the two are the same, so that our invariant image is obtained by simply histogram equalising each of the channels of our original image.

[0026] In the context of image enhancement it is argued [20] that applying an equalisation to the channels of a colour image separately is inappropriate since this can produce significant colour shifts in the transformed image. However our context we are interested not in the visual quality of the image but in obtaining a representation which is illuminant and/or device invariant. Histogram equalisation achieves just this provided that the rank ordering of sensor responses is itself invariant to such changes. In addition, by applying histogram equalisation to each of the colour channels we maximise the entropy in each of those channels. This in itself seems desirable since our intent in computer vision is to use the representation to extract information about the scene and thus maximising the information content of our scene representation ought to be helpful in itself.

4 Invariance of Rank Ordering

[0027] The analysis above shows that the illuminant invariance of histogram equalised images follows directly from the assumption of a diagonal model of illumination change (Equation (2)). But the method does not require Equation (2) to hold to provide invariance. Rather, we require only that rank orderings of sensor responses remain (approximately) invariant under a change in illumination. In fact, the method is not restricted to a change in lighting but to any transformation of the image which leaves rank orderings unchanged. Consider for example Equation (3) which allows the image formation process to include an arbitrary non-linear transform (denoted by $f()$) of sensor responses). Different transforms $f()$, lead of course to different images. But note that, the rank ordering of sensor responses will be preserved provided that $f()$ is a monotonic increasing function. Thus, histogram equalised images are invariant to monotonic increasing functions. This fact is important because many of the transformations such as gamma or tone-curve corrections which are applied to images, satisfy the condition of monotonicity.

[0028] To investigate further the rank invariance of sensor responses we conducted a similar experiment to that of Dannemiller [2] who investigated to what extent the responses of cone cells in the human eye maintain their rank ordering under a change in illumination. He found that to a very good approximation rank orderings were maintained. Here, we extend the analysis to investigate a range of different devices in addition to a range of illuminants. To investigate the invariance of rank orderings of sensor responses for a single device under changing illumination we proceed as follows. Let R_k represent the spectral sensitivity of the k th sensor of the device we wish to investigate. Now suppose we calculate (according to Equation (1)) the responses of this sensor to a set of surface reflectance functions under a fixed illuminant $E^1(\lambda)$. We denote those responses by the vector P^1_k . Similarly we denote by P^2_k the responses of the same sensor to the same surfaces viewed under a second illuminant $E^2(\lambda)$. Next we define a function $rank()$ which takes a vector argument and returns a vector whose elements contain the rank of the corresponding element in the argument.

[0029] Then, if sensor responses are invariant to the illuminants E^1 and E^2 , the following relationship must hold:

$$rank(P^1_k) = rank(P^2_k) \quad (11)$$

[0030] In practice the relationship in Equation (11) will hold only approximately and we can assess how well the relationship holds using Spearman's Rank Correlation coefficient which is given by:

$$\rho = 1 - 6 \sum_{j=1}^N \frac{d_j^2}{N_s(N_s^2 - 1)} \quad (12)$$

where d_j is the difference between the j th elements of $\text{rank}(P_k^1)$ and $\text{rank}(P_k^2)$ and N_s is the number of surfaces. This coefficient takes a value between 0 and 1: a coefficient of zero implies that Equation (11) holds not at all, while a value of one will be obtained when the relationship is exact. Invariance of rank ordering across devices can be assessed in a similar way by defining two vectors: P_k^1 defined as above and Q_k representing sensor responses of the k th class of sensor of a second device under the illuminant E^1 . By substituting these vectors in Equation (12) we can measure the degree of rank correlation. Finally we can investigate rank order invariance across device and illumination together by comparing, for example, the vectors P_k^2 and Q_k^1 .

[0031] We conducted such an analysis for a variety of imaging devices and illuminants, taking as our surfaces, a set of 462 Munsell chips [21] which represent a wide range of reflectances that might occur in the world. For illuminants we chose 16 different lights, including a range of daylight illuminants, Planckian blackbody radiators, and fluorescent lights, again representing a range of lights which we will meet in the world. Finally, we used the spectral sensitivities of the human colour matching functions [21] as well as those of four digital still cameras and a flatbed scanner.

Table 1: Spearman's Rank Correlation Coefficient for each sensor of a range of devices. Results are averaged over all pairs of a set of 16 illuminants.

	1 st Sensor	2 nd Sensor	3 rd Sensor
Across Illumination			
Colour Matching Functions	0.9957	0.9922	0.9992
Camera 1	0.9983	0.9984	0.9974
Camera 2	0.9978	0.9938	0.9933
Camera 3	0.9979	0.9984	0.9972
Camera 4	0.9981	0.9991	0.9994
Scanner	0.9975	0.9989	0.9995
Across Devices			
	1 st Sensor	2 nd Sensor	3 rd Sensor
Illuminant			
Daylight (D65)	0.9877	0.9934	0.9831
Fluorescent (cwf)	0.9931	0.9900	0.9710
Tungsten (A)	0.9936	0.9814	0.9640
Across Device and Illuminant			
	0.9901	0.9886	0.9774

[0032] Table 1 summarises the results for the case in which sensor is fixed and illumination is allowed to change, using the measure defined by Equation (12). Values are shown for each device averaged over all 16 illuminants in our illuminant set, for three illuminants (daylight, fluorescent, and tungsten) averaged over all devices, and for all devices and illuminants. In all cases, the results show a very high degree of correlation: average correlation never falls below 0.964 which represents a high degree of correlation. Minimum correlation over all devices and illuminants was 0.9303 for the 1st sensor, 0.9206 for the 2nd sensor and 0.8525 for the 3rd. Thus on the basis of these results we conclude that rank orderings are preserved to a very good approximation across a change in either or both, device and illumination.

5 An Application to Colour Indexing

[0033] To test the invariance properties of histogram equalisation further we applied the method to an image retrieval task. Finlayson *et al* [3] recently investigated whether existing invariant approaches were able to facilitate good enough

image indexing across a change in either, or both illumination and device. Their results suggested that the answer to this question was no. Here we repeat their experiment but using histogram equalised images as our basis for indexing to investigate what improvement, if any, the method brings.

Table 2: Average Match Percentile results of the indexing experiment for four different cases: (1) Across Illumination, (2) Across cameras, (3) Across scanners, and (4) Across all devices.

Colour Model	Case (1)	Case (2)	Case (3)	Case (4)
RGB	63.23	71.85	98.88	65.53
Greyworld	93.96	94.22	99.34	92.28
Hist Eq.	96.72	95.52	98.94	94.54

[0034] The experiment is based on a database of images of coloured textures captured under a range of illuminants and devices and described in [4]. In summary there are 28 different coloured textures each captured under six different devices (4 cameras and 2 scanners). In addition each camera was used to capture each of the textures under 3 different lights so that in total there are $(3 \times 4 + 2) \times 28 = 392$ images. In image indexing terms this is a relatively small database and it is chosen because it allows us to investigate performance across a change in illumination and device. In our experiments we tested indexing performance across three different conditions: **(1)** across illumination, **(2)** across homogenous devices, and **(3)** across heterogeneous devices **(4)**. In each case the experimental procedure was as follows. First, we choose a set of 28 images all captured under the same conditions (same device and illuminant) to be our image database. Next we select from the remaining set of images a subset of appropriate query images. So, if we are testing performance across illumination, we select as our query images the 56 images captured by the device corresponding to the database images, under the two non-database illuminants. Then for all database and query images we derive an invariant image using either the histogram equalisation method set forth above, or one of a range of previously published [9, 5, 6, 10, 18] invariant methods. Finally we represent the invariant image by its colour distribution: that is, by a histogram of the pixel values in the invariant image. All results reported here are based on 3-dimensional histograms of dimension $16 \times 16 \times 16$.

[0035] Indexing is performed for a query image by comparing its histogram to each of the histograms of the database images. The database image whose histogram is most similar to the query histogram is retrieved as a match to the query image. We compare histograms using the intersection method described by Swain *et al* [19] which we found to give the best results on average. Indexing performance is measured using average match percentile [19] which gives a value between 0 and 100%. A value of 99% implies that the correct image is ranked amongst the top 1% of images whilst a value of 50% corresponds to the performance we would achieve using random matching.

[0036] Table 2 summarises the average match percentile results for the four different conditions. In addition to results for histogram equalisation we also show results based on histograms of the original images (*RGB*), and on Greyworld normalised images, that is on images calculated according to Equation (3). Results for a variety of other invariant representations can be found in [3]: all perform significantly worse than Greyworld. Significantly histogram equalisation outperforms Greyworld for all conditions. Histogram equalisation results across a change in illumination are very good: an AMP of close to 97% as compared to 94% for the second best method. Results of matching across homogeneous devices (Cases 2 and 3 in Table 2) show that both Greyworld and histogram equalisation perform similarly with histogram equalisation performing slightly better on average. Finally across heterogeneous devices histogram equalisation performs best.

[0037] While the results are quite good: the histogram equalisation method clearly outperforms all other methods on average, the experiment does raise a number of issues. First, it is surprising that one of the simplest invariants -- Greyworld-performs as well as it does, whilst other more sophisticated invariants perform very poorly. This performance indicates that for this dataset a diagonal scaling of sensor responses accounts for most of the change that occurs when illuminant or device is changed. It also suggests that any non-linear transform applied to the device responses post-capture (the function $f()$ in Equation (2)) must be very similar for all devices: most likely a simple power function is applied. Secondly, we might have expected that histogram equalisation would have performed somewhat better than it does. In particular, while illumination invariance is very good, device invariance is somewhat less than we might have hoped for given the analysis in Section 4. An investigation of images for which indexing performance was poor reveals a number of artefacts of the capture process which might account for the performance. First, a number of images captured under tungsten illumination have values of zero in the blue channel for many pixels. Second, a number of the textures have uniform backgrounds but the scanning process introduces significant non-uniformities in these regions. For both cases the resulting histogram equalised images are far from invariant. Excluding these images leads to a significant improvement in indexing performance. However, for an invariant image representation to be of practical use in an uncalibrated envi-

ronment it must be robust to the limitations of the imaging process. Thus we have reported results including all images. And we stress again in summary, that the simple technique of histogram equalisation, posited only on the invariance of rank ordering across illumination and/or device outperforms all previous invariant methods and in particular gives excellent performance across changes in illumination.

References

[0038]

[1] J. Bach, C. Fuller, A. Gupta, A. Hampapur, B. Horowitz, R. Humphrey, and R. Jain. The virage image search engine: An open framework for image management. In SPIE Conf. on Storage and Retrieval for Image and Video Databases, volume 2670, pages 76-87, 1996.

[2] James L. Dannemiller. Rank orderings of photoreceptor photon catches from natural objects are nearly illuminant-invariant. *Vision Research*, 33 (1) : 131- 140, 1993.

[3] G. Finlayson and G. Schaefer. Colour indexing across devices and viewing conditions. In 2nd International Workshop on Content-Based Multimedia Indexing, 2001.

[4] G. Finlayson, G. Schaefer, and G. Y. Tian. The UEA uncalibrated colour image database. Technical Report SYS-COO- 07, School of Information Systems, University of East Anglia, Norwich, United Kingdom, 2000.

[5] G.D. Finlayson, S.S. Chatterjee, and B.V. Funt. Color angular indexing. In The Fourth European Conference on Computer Vision (Vol II), pages 16-27. European Vision Society, 1996.

[6] G.D. Finlayson, B. Schiele, and J. Crowley. Comprehensive colour image normalization. In *eccv98*, pages 475-490, 1998.

[7] Graham Finlayson. Coefficient Colour Constancy. PhD thesis, Simon Fraser University, 1995.

[8] Brian Funt, Kobus Baxnard, and Lindsay Martin. Is machine colour constancy good enough? In 5th European Conference on Computer Vision, pages 455-459. Springer, June 1998.

[9] Brian V. Funt and Graham D. Finlayson. Color Constant Color Indexing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17 (5) : 522- 529, 1995.

[10] T. Gevers and A.W.M. Smeulders. Color based object recognition. *Pattern Recognition*, 32:453-464, 1999.

[11] Graham. D. Finlayson, Steven Hordley, and Paul Hubel. Illuminant estimation for object recognition. *COLOR research and application*, 2002. to appear.

[12] G. Healey. Using color for geometry-insensitive segmentation. *Journal of the Optical Society of America, A.*, 6: 920-937, 1989.

[13] B.A. Maxwell and S.A. Shafer. A framework for segmentation using physical models of image formation. In *Computer vision and Pattern recognition*, pages 361-368. IEEE, 1994.

[14] W. Niblack and R. Barber. The QBIC project: Querying images by content using color, texture and shape. In *Storage and Retrieval for Image and Video Databases I*, volume 1908 of SPIE Proceedings Series. 1993.

[15] C. Poynton. The rehabilitation of gamma. In SPIE Conf. Human Vision and Electronic Imaging III, volume 3299, pages 232-249, 1998.

[16] Yogesh Raja, Stephen J. McKenna, and Shaogang Gong. Colour model selection and adaptation in dynamic scenes. In 5th European Conference on Computer Vision, pages 460-474. Springer, June 1998.

[17] B. Schiele and A. Waibel. Gaze tracking based on face-color. In *International Workshop on Automatic Face- and Gesture-Recognition*, June 1995.

[18] M. Sticker and M. Orengo. Similarity of color images. In SPIE Conf. on Storage and Retrieval for Image and Video Databases III, volume 2420, pages 381-392, 1995.

[19] Michael J. Swain and Dana H. Ballard. Color Indexing. International Journal of Computer Vision, 7 (1) : 11- 32, 1991.

[20] Alan Watt and Fabio Policarpo. The Computer Image. Addison-Wesley, 1997.

[21] G. Wyszecki and W.S. Stiles. Color Science: Concepts and Methods, Quantitative Data and Formulas. New York:Wiley, 2nd edition, 1982.

Claims

1. Method for image retrieval, comprising:

representing an original colour image having values in a plurality of colour channels for each of a plurality of points within the original image by a transformed image obtained by histogram equalizing each of the colour channels of the original colour image; and

comparing the resulting histogram of the transformed image to the respective histograms of each image of a set of similarly transformed colour images, wherein the image in the set of colour images whose histogram is most similar to the histogram of the transformed image is retrieved as a match to the original colour image.

2. Method according to claim 1, comprising the step of representing the transformed image by its colour distribution.

3. Method according to claim 2, in which the transformed image and each of the similarly transformed colour images in the set of colour images is represented by a histogram of pixel values in the respective transformed image.

4. Method according to claim 1, further comprising the step of normalizing each histogram.

5. Method according to any preceding claim, in which each point within the original image is a pixel;

6. Method according to any preceding claim in which the colour channels are RGB channels.

7. Method according to any preceding claim, in which the plurality of points within the original image represents a part of the original image;

8. Method according to claim 7, wherein the part of the original image comprising proximate points or pixels.

9. An image processing apparatus for carrying out the method as defined in any preceding claim.

10. A computer programmed to carry out the method as defined in any of claims 1 to 8.

11. A computer-readable medium carrying a program for causing a general purpose computer to carry out the method as defined in any of claims 1 to 8.

Patentansprüche

1. Verfahren zum Abrufen von Bildern, aufweisend:

Darstellen eines Originalfarbbildes, welches für jeden einer Vielzahl von Punkten innerhalb des Originalbildes Werte für eine Vielzahl von Farbkanälen aufweist, durch ein transformiertes Bild, welches durch eine Histogrammäqualisation von jedem der Farbkanäle des Originalfarbbildes erhalten wird; und

Vergleichen des erhaltenen Histogramms des transformierten Bildes mit den jeweiligen Histogrammen von jedem Bild aus einer Menge von ähnlich transformierten Farbbildern, wobei das Bild aus der Menge der Farbbilder, dessen Histogramm dem Histogramm des transformierten Bildes am ähnlichsten ist, als ein zu dem Originalbild passendes Bild abgerufen wird.

2. Verfahren nach Anspruch 1, aufweisend den Schritt des Darstellens des transformierten Bildes durch seine Farbverteilung.
3. Verfahren nach Anspruch 2, bei welchem das transformierte Bild und jedes der ähnlich transformierten Farbbilder aus der Menge der Farbbilder durch ein Histogramm von Pixelwerten in dem jeweiligen transformierten Bild dargestellt werden.
4. Verfahren nach Anspruch 1, weiterhin aufweisend den Schritt des Normalisierens jedes Histogramms.
5. Verfahren nach irgendeinem der vorhergehenden Ansprüche, bei welchem jeder Punkt innerhalb des Originalbildes ein Pixel darstellt.
6. Verfahren nach irgendeinem der vorhergehenden Ansprüche, bei welchem die Farbkanäle RGB-Kanäle darstellen.
7. Verfahren nach irgendeinem der vorhergehenden Ansprüche, bei welchem die Vielzahl von Punkten innerhalb des Originalbildes einen Teil des Originalbildes darstellt.
8. Verfahren nach Anspruch 7, wobei der Teil des Originalbildes benachbarte Punkte oder Pixel aufweist.
9. Bildverarbeitungsvorrichtung zum Ausführen des Verfahrens, wie in irgendeinem vorhergehenden Anspruch definiert.
10. Computer, der programmiert ist, das Verfahren nach irgendeinem der Ansprüche 1 bis 8 auszuführen.
11. Computerlesbares Medium zum Tragen eines Programms, welches einen Universalcomputer dazu veranlasst, das Verfahren, wie in einem der Ansprüche 1 bis 8 definiert, auszuführen.

Revendications

1. Procédé pour retrouver une image, comprenant:

la représentation d'une image d'origine en couleurs possédant des valeurs dans une pluralité de canaux de couleur pour chacun d'une pluralité de points au sein de l'image d'origine par une image transformée obtenue par un histogramme égalisant chacun des canaux de couleur de l'image d'origine en couleurs; et
la comparaison de l'histogramme résultant de l'image transformée aux histogrammes respectifs de chaque image d'un jeu d'images transformées en couleurs de façon semblable, l'image du jeu d'images en couleurs dont l'histogramme est le plus semblable à l'histogramme de l'image transformée étant retrouvée comme concordant avec l'image d'origine en couleurs.
2. Procédé selon la revendication 1, comprenant l'étape de représentation de l'image transformée par sa distribution de couleurs.
3. Procédé selon la revendication 2, dans lequel l'image transformée et chacune des images transformées en couleurs de façon semblable du jeu d'images en couleurs sont représentées par un histogramme de valeurs de pixel dans l'image transformée respective.
4. Procédé selon la revendication 1, comprenant en outre l'étape de normalisation de chaque histogramme.
5. Procédé selon l'une des revendications précédentes, dans lequel chaque point au sein de l'image d'origine est un pixel.
6. Procédé selon l'une des revendications précédentes, dans lequel les canaux de couleur sont des canaux RVB.
7. Procédé selon l'une des revendications précédentes, dans lequel la pluralité de points au sein de l'image d'origine représente une partie de l'image d'origine.
8. Procédé selon la revendication 7, dans lequel la partie de l'image d'origine comprend des pixels ou des points

proches.

9. Un dispositif de traitement d'image pour mettre en oeuvre le procédé défini dans l'une des revendications précédentes.

5

10. Un calculateur programmé pour mettre en oeuvre le procédé défini dans l'une des revendications 1 à 8.

11. Un support lisible par calculateur supportant un programme pour faire en sorte qu'un calculateur universel mette en oeuvre le procédé défini dans l'une des revendications 1 à 8.

10

15

20

25

30

35

40

45

50

55

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Non-patent literature cited in the description

- **CHEN et al.** Rotation and gray scale transform invariant texture identification using wavelet decomposition and hidden Markov model. *Transactions on Pattern Analysis and Machine Intelligence*, February 1994, vol. 16 [0011]
- **J. BACH ; C. FULLER ; A. GUPTA ; A. HAMPAPUR ; B. HOROWITZ ; R. HUMPHREY ; R. JAIN.** The virage image search engine: An open framework for image management. *SPIE Conf. on Storage and Retrieval for Image and Video Databases*, 1996, vol. 2670, 76-87 [0038]
- **JAMES L. ; DANNEMILLER.** Rank orderings of photoreceptor photon catches from natural objects are nearly illuminant-invariant. *Vision Research*, 1993, vol. 33 (1), 131-140 [0038]
- **G. FINLAYSON ; G. SCHAEFER.** Colour indexing across devices and viewing conditions. *2nd International Workshop on Content-Based Multimedia Indexing*, 2001 [0038]
- **G. FINLAYSON ; G. SCHAEFER ; G. Y. TIAN.** The UEA uncalibrated colour image database. *Technical Report SYS-COO-07, School of Information Systems, University of East Anglia, Norwich, United Kingdom*, 2000 [0038]
- Color angular indexing. **G.D. FINLAYSON ; S.S. CHATTERJEE ; B.V. FUNT.** The Fourth European Conference on Computer Vision. European Vision Society, 1996, vol. II, 16-27 [0038]
- **G.D. FINLAYSON ; B. SCHIELE ; J. CROWLEY.** Comprehensive colour image normalization. *eccv98*, 1998, 475-490 [0038]
- Coefficient Colour Constancy. **GRAHAM FINLAYSON.** PhD thesis. Simon Fraser University, 1995 [0038]
- Is machine colour constancy good enough?. **BRIAN FUNT ; KOBUS BAXNARD ; LINDSAY MARTIN.** 5th European Conference on Computer Vision. Springer, June 1998, 455-459 [0038]
- **BRIAN V. FUNT ; GRAHAM D. FINLAYSON.** Color Constant Color Indexing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1995, vol. 17 (5), 522-529 [0038]
- **T. GEVERS ; A.W.M. SMEULDERS.** Color based object recognition. *Pattern Recognition*, 1999, vol. 32, 453-464 [0038]
- **GRAHAM. D. FINLAYSON ; STEVEN HORDLEY ; PAUL HUBEL.** Illuminant estimation for object recognition. *COLOR research and application*, 2002 [0038]
- **G. HEALEY.** Using color for geometry-insensitive segmentation. *Journal of the Optical Society of America, A.*, 1989, vol. 6, 920-937 [0038]
- **B.A. MAXWELL ; S.A. SHAFER.** A framework for segmentation using physical models of image formation. *Computer vision and Pattern recognition*, 1994, 361-368 [0038]
- **W. NIBLACK ; R. BARBER.** The QBIC project: Querying images by content using color, texture and shape. *Storage and Retrieval for Image and Video Databases I, volume 1908 of SPIE Proceedings Series*, 1993, vol. 1908 [0038]
- **C. POYNTON.** The rehabilitation of gamma. *SPIE Conf. Human Vision and Electronic Imaging III*, 1998, vol. 3299, 232-249 [0038]
- Colour model selection and adaptation in dynamic scenes. **YOGESH RAJA ; STEPHEN J. MCKENNA ; SHAOANG GONG.** 5th European Conference on Computer Vision. Springer, June 1998, 460-474 [0038]
- **B. SCHIELE ; A. WAIBEL.** Gaze tracking based on face-color. *International Workshop on Automatic Face- and Gesture-Recognition*, June 1995 [0038]
- **M. STICKER ; M. ORENGO.** Similarity of color images. *SPIE Conf. on Storage and Retrieval for Image and Video Databases III*, vol. 2420, 381-392 [0038]
- **MICHAEL J. SWAIN ; DANA H. BALLARD.** Color Indexing. *International Journal of Computer Vision*, 1991, vol. 7 (1), 11-32 [0038]
- **ALAN WATT ; FABIO POLICARPO.** The Computer Image. Addison-Wesley, 1997 [0038]
- **G. WYSZECKI ; W.S. STILES.** Color Science: Concepts and Methods, Quantitative Data and Formulas. New York:Wiley, 1982 [0038]